Talking with Tom Rosamilia...
OS/390: Migration flexibility

Whether you are still running OS/390® V2R6 or have recently moved to a newer release of OS/390, flexible migration paths are in place to allow you to quickly get current. On your behalf, we have made a special exception to our coexistence policy allowing OS/390 V2R6 through z/OS 1.1 to coexist.

OS/390 V2R10 extends the value of OS/390 to include expanded platform choices, improved system ease-of-use, enhanced deployment capabilities and greater product integration to meet tomorrow's e-business needs. The application sourcing models supported include OS/390 traditional batch/OLTP, OS/390 Unix® System Services, LINUX™ for S/390, Java, Enterprise Java Beans (EJB) and CORBA. These models enable the rapid development and deployment of applications and integrate in with mission-critical legacy data. This combination enables the delivery of high-performance, scalable and secure e-business transactions.

OS/390 V2R10 is an excellent stepping stone to z/OS and when combined with the new z/900 server provides the path for exploitation of the new z/Architecture. The z900 is upgradeable from G5/G6, provides up to 50% additional capacity and provides a platform for workload consolidation. z/OS, OS/390 V2R10 and z/VM all exploit 64-bit real
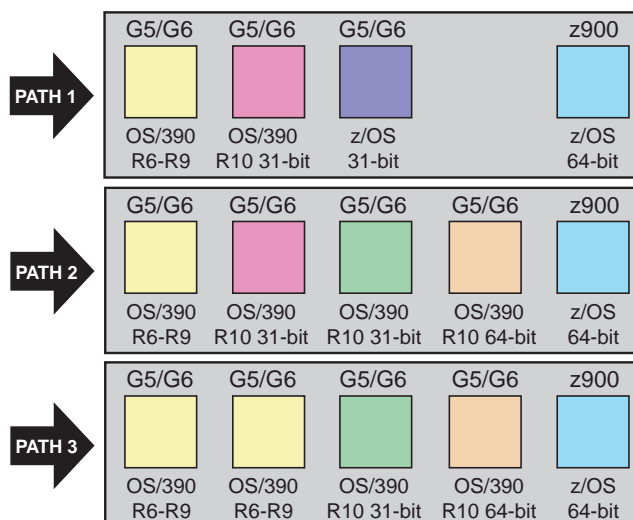
storage removing memory-related bottlenecks due to the ever increasing processor speed and capacity. Linux for zSeries supports both 64-bit real and virtual at this time. Operating system images running in 31-bit and 64-bit can coexist in LPARs side-by-side on the same z900. This coexistence provides a facility for simple testing and migration to 64-bit.

There are multiple possible migration paths to get to the z900 running z/OS in 64-bit mode. You can choose to upgrade either your hardware or your operating system first. It is also possible to move from prior levels of OS/390 directly to z/OS bypassing OS/390 V2R10. The OS/390 V2R10 to z/OS upgrade can also be accomplished by applying the Product Upgrade Package (PUP) to OS/390 V2R10. This procedure is much simpler and faster compared to a complete install. As with OS/390,

you can introduce z/OS in your environment one image at a time, avoiding a multisystem IPL.

The coexistence provided for OS/390, non 64-bit and non-Goal mode images is part of the total migration flexibility provided with all of the zSeries features. New function can be phased-in gradually and major changes are not required to exploit the new function. I recommend that you build a plan to move to WLM Goal mode, upgrade your hardware to the z900 processor, upgrade your operating system to z/OS and Linux for zSeries in the near future. The strengths of the zSeries platform will leverage your current investments and form a technical infrastructure that will launch your business into the next generation of e-business. •

Tom Rosamilia is Vice President, zSeries Software Development

PATH 1

G5/G6	G5/G6	G5/G6		z900
OS/390 R6-R9	OS/390 R10 31-bit	z/OS 31-bit		z/OS 64-bit

PATH 2

G5/G6	G5/G6	G5/G6	G5/G6	z900
OS/390 R6-R9	OS/390 R10 31-bit	OS/390 R10 31-bit	OS/390 R10 64-bit	z/OS 64-bit

PATH 3

G5/G6	G5/G6	G5/G6	G5/G6	z900
OS/390 R6-R9	OS/390 R6-R9	OS/390 R10 31-bit	OS/390 R10 64-bit	z/OS 64-bit

Contents

Enterprise Extender for large network support

RICHARD J. TOBACCO AND BALA RAJARAMAN

Data Center Performance

Are you concerned about the impact to S/390 as you consider SNA and TCP/IP transport network consolidation with Enterprise Extender (EE)? Don't worry, we have collected the 'impact assessment' data that shows even the largest networks will find EE a cost effective solution.

EE is an extension to High Performance Routing, a lower overhead less-complex SNA, that efficiently encapsulates SNA application traffic within UDP frames by EE-capable devices at the edges of an IP network. To the IP network the SNA traffic is UDP data-grams that get routed without hardware or software changes to the IP backbone. To the user the session is 'normal SNA' with predictable performance and high availability. There is no protocol transformation in the network gateway devices into S/390 (such as there is with DLSw) and the encapsulation is performed at the routing layers enabling efficient use of the IP infrastructure for accessing SNA-based data.

EE is available on Communications Server for OS/390 V2R6 (or later) as a base component of OS/390. EE implemented within a branch office traffic consolidation device, as exemplified by IBM Communications Server for Windows, and terminated within System/390 eliminates the need for multiple outboard processors within the data center - reducing management complexity.

With subarea SNA boundary functions are not controlled within the S/390. An outboard device - typically a 3746 Communications Controller, provides message formatting and transmission control. With HPR, end stations ensure data integrity and intermediate devices forward traffic. The intermediate - Automatic Network Routing (ANR) - nodes forward packets, has no session awareness, and relies on the endpoints for error recovery. The endpoint - Rapid Transport Protocol (RTP) - nodes provide end-to-end error recovery, non-disruptive rerouting and selective retransmission of lost packets.

How many connection pipes (routes) can S/390, acting as an RTP endpoint, support? A lot. In this test there were up to 6000 connections or 'pipes'.

Although the number of users and applications that can share a HPR pipe depend on multiple factors, experience suggests that an average of three or more end-users typically share a pipe. Most companies will have significantly less connection pipes and less S/390 burden.

A measure of scalability that permits comparisons between different size processors is CPU-seconds. A CPU-second is the measure of workload capacity the processor(s) perform in a second. A transaction occurring every second might consume 1/25 of the CPU on one processor, but only 1/100 CPU-seconds on another. Refer to the Large Scale Processor Reference (LSPR) material at **http://www.s390.ibm.com/lspr** to relate your processor capacity to the processor used in these tests.

The additional workload of maintaining HPR connections on the S/390 results in a minor increase in overall workload (as compared to subarea SNA). Testing[1] showed EE support across an OSA Express (Gigabit Ethernet) adapter to be more efficient than 'traditional' HPR across an ESCON® channel - throughput was greater and CPU usage was less. However, the additional cost associated with any of the solutions is not significant and should not be a significant part of your network transport decision.

A somewhat easier way of determining the relationship between connections (pipes), transactions, and connectivity option is shown with the '10% CPU Utilization' graph. The number of transactions (per second) supported, when 10% if the CPU is dedicated to transaction processing, is graphed against the number of supported connections (pipes). Although 10% may be a high estimate of the total
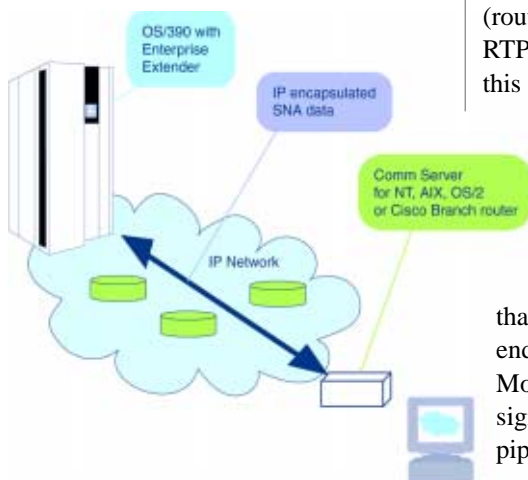
CPU resource typically used on network connectivity, CPU usage scales linearly with transaction rates.

No matter which way you chose to calculate the impact of supporting network connectivity, you will quickly determine that this support is a small burden to most S/390 servers.

Branch office performance

You decide that consolidating your IP and SNA transport networks makes good business sense and at a 'cost' to S/390 that is affordable. Now you must decide what EE-capable devices are available for the remote branch office. Are these devices powerful enough to support the expected traffic? Is EE in a branch office device as effective as the DLSw alternative?

Tests conducted in IBM facilities indicate that IBM Communications Server products (e.g., CS/NT) and a Cisco 3640 with SNASw, their name for the Enterprise Extender function, provide equivalent performance. Furthermore the branch office router provided greater throughput and faster response supporting EE, then when providing DLSw support.

This 'branch office' test consisted of 16 PCs on a LAN, connected to the 'data center' with a T1 link. The EE test configurations were identical with only the EE end-point being changed. EE transaction rates became constrained by the capacity of the link. The DLSw testing was performed with the addition of another DLSw-capable router in front of the S/390. Because DLSw is more process intensive, link capacity was never reached.

Summary

EE is IBM's strategic solution for transporting SNA traffic, end-to-end across TCP/IP networks. Unlike DLSw, EE end-nodes can be situated within S/390 servers. The above performance studies show that running EE end-to-end, with
S/390 as the endpoint does not impact performance or use excessive S/390 resources. If anything, EE has been shown to be extremely affordable and as such to be a very compelling solution for today's larger networks. •

1. All the results described here were obtained in a dedicated test environment. Actual customer results may vary from those shown here. Contact IBM for further information.

S/390 CPU Usage

10% CPU Utilization

Link Utilization

Transaction Rates

Response Times
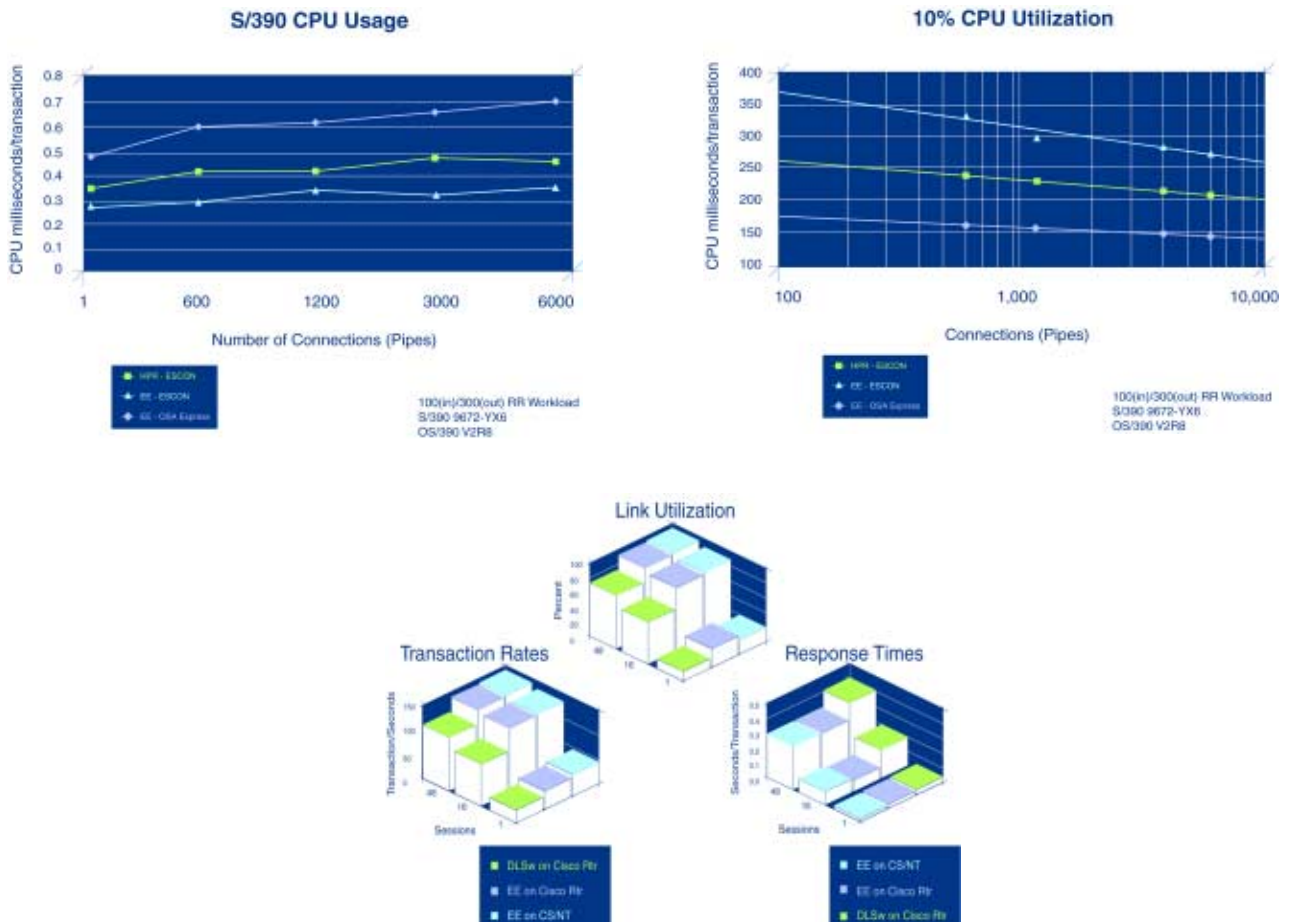
Estimating TN3270 performance

DAN PATEL, RICHARD TOBACCO

How to access those existing SNA applications residing on S/390?

Today many businesses use TN3270 as their method of accessing those SNA applications. The end-users' TCP/IP-based processor emulates a SNA device– they are the TN3270 clients and they talk to TN3270 servers across the IP network. Servers can be outboard (typically a separate channel-attached device) or contained within the S/390.

Does S/390 have the capacity to handle the large number of sessions in your network?

Today's OS/390 telnet server has been stress tested at 64,000 sessions (V2R7 and higher). Architecturally the server can support 64,000 sessions on up to 255 ports – a really big number!

What is the 'cost' of support – how much of my S/390 CPU will be used supporting these telnet sessions?

This depends on which release your S/390 is on – significant improvements in the TCP/IP stack has nearly halved the CPU usage from the early stand-alone TCP/IP versions.

Performance testing, by necessity, is performed in a laboratory using equipment to simulate the infrastructure. An 'Echo' application responds (with 800 bytes) to each (100 byte) request. Because the Echo application invokes VTAM®, the impact on CPU usage is significant but not easily removed. (Customers often estimate the Telnet portion of the measured usage to be 50%.) Measurements are taken using CPU seconds – enabling a mea-surement taking with one proces-sor type to be compared (using LSPR) to another.

In this test[1] example, OS/390 V2R8 used 0.65 milliseconds per transaction. A thousand users entering six transactions per minute or six thousand users entering one transaction per minute would create 100 transactions per second. This transaction rate uses 0.065 CPU seconds or 6.5% of the S/390 9672-RX6. Because customer applications use more CPU than the test application, customers should expect the impact of Telnet on total CPU use to be small.
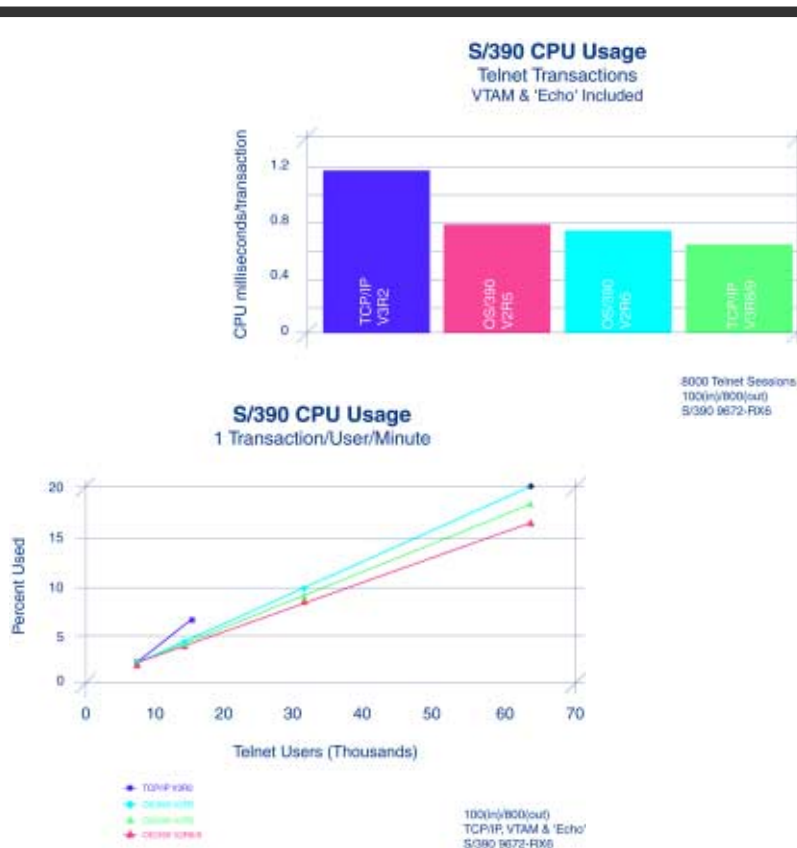
Additional testing has been performed to determine the CPU 'cost' based on the number of sessions. In the following example the transaction rate was held constant at a single transaction per user per minute. Again, the test results include VTAM and the Echo test application.

Summary

Performance improvements have made TN3270 on S/390 affordable – on a par, if not less, than the alternative out-board server solution. And, the additional cost of a TN3270 server on S/390 is generally only a small compo-nent part of the total application. Scalability of the TN3270 server on OS/390 surpasses the require-ments of most companies. And, the inboard server setup and management is considerably easier than out-board solutions – but that's a topic of another article. •

1. All the results described here were obtained in a dedicated test environment. Actual customer results may vary from those shown here. Contact IBM for further information.

S/390 CPU Usage
Telnet Transactions
VTAM & 'Echo' Included

8000 Telnet Sessions
100(in)/(800(out)
S/390 9672-RX6

S/390 CPU Usage
1 Transaction/User/Minute

100(in)/800(out)
TCP/IP, VTAM & 'Echo'
S/390 9672-RX6

Order S/390 Software - as simple as a few "clicks"

ROBERT TERWILLIGER

Planning and ordering software upgrades can be a time consuming and complex task. Just trying to figure out what to order to upgrade your current system can require lots of research. Are your tired of sorting through RFAs, product checklists, etc. trying to figure out product migration paths? Now there's a fast, easy way to order your S/390 software.

Kerry Collia

Just go to ShopS390 at **http:// inetsd01.boulder.ibm.com/ ShopS390/ShopS390.jsp** and upload your installed inventory from your target system. Then, select the desired upgrade offering such as a ServerPac or CBPDO.

ShopS390 automatically analyzes your installed inventory and maps it to your IBM licensed inventory records and the latest IBM software product catalog. It gives you a preselection of your upgrade options. Then, it lets you perform technical requisite checking and various "what ifs." If you choose a release upgrade that requires no new software license, you can submit your order on the spot and track the order status through delivery.

However, if your order contains a "priced" upgrade and/or requires new software licenses, it will be routed through the appropriate IBM representative. You'll receive an e-mail confirmation of your order, and you'll be contacted by an IBM representative who will give you a price quote. Once you accept, the sales representative will submit your order for processing. You order is saved on ShopS390 for your future reference.

Soon, you'll have the option of viewing just your licensed inventory without also viewing your installed inventory and the latest available products. You'll also be able to view version upgrade opportunities across your entire establishment.

ShopS390 allows corporations to manage access and level of authority within their own enterprise. Plus, it allows customers to give Business Partners authority to place orders on their behalf via ShopS390. Go to ShopS390 today at **http:// inetsd01.boulder.ibm.com/ ShopS390/ShopS390.jsp** to learn more about this new and exciting e-business solution.

Please note:
ShopS390 is currently only available in the U.S. It will be available in EMEA later this fall. Availability in AP and Canada is targeted for the first half of 2001. Please stayed tuned to the ShopS390 Web site for availability in your locale. •

64-bit real operation in z/OS

JIM GROUND

z/OS operates in the new 64-bit z/Architecture on a zSeries processor. (z/OS operates in the older 31-bit ESA/390 architecture on an S/390 G5 or G6 processor.) While OS/390 Release 10 gives you the choice to run in either architecture mode on a zSeries machine, z/OS must run under the newer z/Architecture on a zSeries processor. The benefits of 64-bit operation are the same for z/OS Release 1 as they are for OS/390 Release 10.

For more information about the benefits of 64-bit real operation, see *64-bit real operation in OS/390 Release 10* in the z/OS section of this newsletter.

Another (!) New Redbook!

ANDREW D. SCHMIDT

In the August 2000 issue of "Hot Topics," I wrote an article discussing an International Technical Support Organization (ITSO) Redbook titled *"Experiences Moving a Java Application to OS/390 , SG24-5620"* that was published in late 1999. The project that produced that Redbook was funded by my team, PartnerWorld for Developers, S/390.

PartnerWorld for Developers is a worldwide program supporting developers who build solutions using IBM technologies, and the S/390 (and now zSeries) team that supports those ISVs is located at the home of OS/390 in Poughkeepsie, New York. As a Redbook primarily addressing a Java application port, it contains much useful information about Java for OS/390 that is of general interest to OS/390 Java programmers/developers, regardless of whether a port was being done or a new application was being provided for the S/390 platform.

Java, of course, is one hot language for S/390 running under OS/390's domain. Another is C/C++. I'm proud to say that my team did it again in 2000! We just finished producing another Redbook (to be published in late fall 2000, but already available on the web **http://www.redbooks.ibm.com** as a Redpiece) entitled *"C/C++ Applications on OS/390 UNIX, SG24-5992."* This project was lead by Michael MacIssac of the ITSO in the summer of 2000, assisted by other IBMers and ISV vendors. The book documents the experiences in moving applications written in C/C++ from other UNIX® platforms to OS/390 UNIX

Kim Mulé

System Services ('OS/390 UNIX'). The book applies to OS/390 Version 2 Releases 9 and 10.

As the preface states, the book highlights the traditional strengths of OS/390 and describes some of the subsystems not always found on other UNIXes. It addresses application development tools, the C/C++ compiler, some open source development code such as the C++ Standard Template Library (STL) and Adaptive Communication Environment (ACE). Also performance techniques are suggested. The first seven chapters of this book describe the C/C++ environment of OS/390 in general.

Like its Java Redbook predecessor, the Redbook not only describes the environment of running C/C++ applications on OS/390, but also describes the real world events relating to the porting of the application itself in terms of the planning and execution activities. Of great interest is tuning of the application, and an entire chapter is devoted to performance and tuning techniques.

A liberal amount of appendixes are provided with sample JCL procedures, code snippets, comparisons of GNU vs. the OS/390 compiler, ASCII support information, and Dumps. Once again, although written from the point of view of a porting application project, this Redbook is of great interest to the general OS/390 application audience that provides apps written in the C/C++ programming language. •

Have you heard about Linux for S/390? This new operating system on the S/390 platform will allow you to quickly and easily port applications to run on your S/390 hardware and greatly increase the size of your application portfolio.

If you would like to know more about running this exciting, fast-growing operating system on your S/390 hardware, see the article "Linux meets the enterprise servers" on the z/OS side of this newsletter.

OS/390 Support for Unicode™

JOACHIM BAUER

The Unicode™ Standard[1]

The Unicode Standard is the universal character encoding standard used for the representation of text for computer processing. Unicode provides a consistent way of encoding multilingual plain text and brings order to a chaotic state of affairs that has made it difficult to exchange text files internationally. Computer users who deal with multilingual text or who regularly use mathematical symbols and other technical characters will find that the Unicode Standard greatly simplifies their work.

Paul Rowntree

The design of Unicode is based on the simplicity and consistency of ASCII, but goes far beyond ASCII's limited ability to encode only the Latin alphabet. The Unicode Standard provides the capacity to encode all of the characters used for the written languages of the world. It uses a default 16-bit encoding that provides code points for more than 65,000 characters. An extension mechanism called UTF-16 allows the encoding of a million more characters by using two 16 bit entities (called surrogate pair) for one character. This is sufficient for all known character encoding requirements, including full coverage of all historic scripts of the world.

For each code point the Unicode Standard provides information about each character in the readable file UnicodeData.txt, such as the general category (letter, separator,...) and the corresponding uppercase, lowercase, and titlecase mappings.

The latest version of the Unicode Standard is 3.0. For further information on the standard, go to the Unicode® Consortium's web site at **http://www.unicode.org/**.

OS/390 V2 R8/R9/R10 Support for Unicode™

OS/390 Support for Unicode brings Unicode to OS/390. Its first release offers character conversion as well as basic case conversion. Within character conversion, characters are converted from one coded character set identifier (CCSID) to another. Basic case conversion allows conversion between upper- and lowercase based solely on the file UnicodeData.txt provided by the Unicode Consortium. It does not include special casing such as Greek 'S' (capital sigma) that lowercases to 'ρ' (small sigma) if it is followed by another letter, but lowercases to 's' (small final sigma) if it is not.

OS/390 Support for Unicode consists of two main components:

- the infrastructure which creates the conversion environment
- the conversion services which use the conversion environment

The conversion services can only be used when the conversion environment is active. The infrastructure provides tools to create a conversion image. When the image is loaded into a common data space, the conversion environment is activated and the conversion services are ready to be used by callers.

The infrastructure

The infrastructure supplies all parts to customize and initialize the conversion services. The services are activated during IPL. All necessary resources are allocated and conversion tables are loaded into storage when defined in the conversion image. Changes to the conversion environment can also be done in a running system. Basically, the purpose of the infrastructure is to establish and maintain the conversion environment.

The user can specify the required conversions and control information in a text file. IBM provides a library of conversion tables. In addition the user can add user-defined conversion tables.

The conversion services

The other major part of OS/390 Support for Unicode are the conversion services. They consist of a variety of conversion types and permit character conversion, as well as case conversion, for which the appropriate conversion tables were provided by the infrastructure.

The conversion services are provided as callable system services and address functions formerly available only in C/C++ and the Language Environment. Some programs have requirements to exploit these functions, but run in

system modes that are incompatible with the Language Environment's run-time environment. The conversion services are provided with an HLASM interface and do not expect many prerequisites concerning the system mode (see the OS/390 V2 R8/R9/R10 Support for Unicode documentation for details).

A "C" interface is also provided for programs that run in the Language Environment.

Character conversion

OS/390 Support for Unicode provides the conversion of a set of characters from one CCSID (coded character set identifier) to another CCSID. A basic feature of a CCSID is its encoding scheme. CCSIDs and encoding schemes are defined in *Character Data Representation Architecture Reference and Registry, SC09-2190.* Figure 1 lists which of the encoding schemes defined in the Charac-

ter Data Representation Architecture (CDRA) are supported.

UCS-2 may be encoded in big endian or in small endian format. OS/390 Support for Unicode, however, only handles the big endian format. OS/390 Support for Unicode does not handle surrogate pairs except for conversions to and from UTF-8.

Code pages with a pure single-byte or pure double-byte encoding (SBCS, DBCS, and UCS-2) are called simple code pages. Code pages that consist of two or more sub-code pages (PC MBCS, EUC MBCS, EBCDIC MBCS, and ISO2022 MBCS) are called mixed code pages.

Conversions between simple code pages are supported directly. Besides these direct conversions, indirect conversions also exist where any CCSID can be converted into another by specifying

an intermediate CCSID. Normally, the intermediate CCSID is a UCS-2 CCSID. The indirect conversion via UCS-2 is used by OS/390 Support for Unicode automatically, if there is no table available for direct conversions between CCSID 1 and CCSID 2, but if there are tables available from CCSID 1 to UCS-2 and from UCS-2 to CCSID 2.

Case conversion

OS/390 Support for Unicode provides basic case conversions that allow the conversion of Unicode characters to their uppercase or lowercase equivalent. For more details about the case mappings, refer to the tables provided by the Unicode Consortium. Currently only the basic case conversion is supported as described in the file UnicodeData.txt at the Unicode Consortium Web site at **http://www.unicode.org/.**

The registration/download Web site for Unicode is **http://www6.software.ibm.com/dl/os390/unicodespt-p.** For more information, see *Support for Unicode: Using Conversion Services SC33-7050-00.* •

	Encoding Scheme	Description
Simple	SBCS	Pure single-byte encoding
	DBCS	Double single-byte encoding
	UCS-2	Universal character set encoded in 2 octets (ISO/IEC 10646 encoding form)
	UTF-8	Unicode Transformation Format, 8-bit encoding form
Mixed	PC MBCS	IBM-PC mixed single byte and double-byte encoding with implicit code extension
	EUC MBCS	IBM EUC encoding
	EBCDIC MBCS	EBCDIC mixed single-byte and double-byte encoding using SI/SO code extension
	ISO2022 MBCS	ISO2022 TCP/IP encoding

Fig. 1: The encoding schemes